Defense Technical Information Center Compilation Part Notice

ADP023863

TITLE: Guide to the Best and Worst HPCMP Systems for Executing Individual TI-04 Benchmarking Applications and Synthetics

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Proceedings of the HPCMP Users Group Conference 2004. DoD High Performance Computing Modernization Program [HPCMP] held in Williamsburg, Virginia on 7-11 June 2004

To order the complete compilation report, use: ADA492363

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP023820 thru ADP023869

Guide to the Best and Worst HPCMP Systems for Executing Individual TI-04 Benchmarking Applications and Synthetics

Roy L. Campbell, Jr.
US Army Research Laboratory (ARL), Computational and Information Sciences Directorate, Aberdeen Proving Ground, MD
rcampbell@arl.army.mil

Abstract

Since time-to-solution and processor scalability of an application can vary greatly from one architecture to another, it is important to consider the suitability of each system with respect to that application in order to make efficient use of available resources. Given that the full set of possible applications is quite large, this paper focuses on those applications within the FY-04 technology insertion (TI-04) benchmarking suite – AERO, COBALT-60, GAMESS, HYCOM, NAMD, OOCore, and RF-CTH – for which all except AERO have a standard and large test case. Both the overall performance and performance per processor of the High Performance Computing Modernization Program's (HPCMP's) major systems are analyzed for each test case in order to provide general selection guidance to users and a unique perspective for code developers.

1. Introduction

Performance variation among architectures is not a new topic, especially to those in the mid 90's responsible for converting vector codes to efficiently use more commodity-like instruction set architectures (ISAs). Now, almost ten years later, vector architectures are back. This time they are just one of a multitude of architectures that are available to users, making the mapping of problems to systems that much more difficult. So, assuming a user has a particular application and maybe even a particular problem in mind, how does he or she decide what system to use? Some may decide to continue to use a system they have used before, but the lifespan of that system is limited (to typically 3.5 years). Some may decide to stay with a particular vendor, but market forces can in some cases cause yesterday's winners to be tomorrow's losers. Some may decide to remain with a general architecture, but the industry is in a continual process of reinventing itself as outside influences such as Government sponsored research encourage vendors to "think outside of the box" and "dream big". So it is worthwhile for a user to re-examine his or her choice on a periodic basis (e.g., once a year), but on what basis? Fortunately, the HPCMP has decided as a matter of policy to assess all of its systems, both new and old, using its annually updated technology insertion (TI-XX) benchmarking suite, with the first comprehensive assessment being performed using the TI-04 suite (i.e., the suite for the most recently completed program acquisition). Therefore, a rich set of data is available to compare existing program systems.

The usefulness of the comparison will lie in how the well the user's application and problem (or test case) maps to a particular application and test case in the benchmarking suite. As a cursory guide, discipline associations for each TI-04 benchmarking code are provided below:

<u>CCM</u>
GAMESS – quantum chemistry
NAMD – molecular dynamics
<u>CEA</u>
OOCORE – electromagnetics
<u>CFD</u>
AERO – aeroelastic fluid/structure interactions
COBALT-60 – general flow (Euler/Navier-Stokes)
<u>CSM</u>
RFCTH – shock physics
<u>CWO</u>
HYCOM – ocean modeling

Additional code descriptions can be found in Tracy et al., 2003[1].

2. Problem and Methodology

Work is often classified into one of two types – capability-oriented and capacity-oriented. In the extreme

case, capability-oriented problems are so large that all processors are required on a well-balanced, state-of-the-art system in order to reduce the time-to-solution to a reasonable fraction of the system's mean-time-to-failure. For such problems, the total capability of a system is of interest and is determined by the time-to-solution for the problem when using all processors. Capacity-oriented problems, on the other hand, have reasonable time-to-solutions on a fraction of the same system, but often require a large number of independent executions to cover a host of scenarios. In that case, the total capacity of the system is of interest and is deduced by determining how many scenarios that system can execute in a given amount of time. For users that require high single image capability or capacity, performance results are provided, while for users that require moderate single image capability and capacity, or require high capacity but are willing to spread work across a number of systems, performance per processor results are provided.

Twenty-seven major unclassified HPCMP systems are included in this study, three of which are systems that were purchased in the TI-04 acquisition, and are therefore identified by non-descriptive monikers (e.g., System A). Actual system identifications can be provided along with verified/updated results, once installation by the respective vendors and acceptance/benchmarking by the Government for these systems has been completed.

3. Results

For each application test case, the major unclassified HPCMP systems were ranked by performance (Table 1) and performance per processor (Table 2) with the top and bottom five systems denoted in tan and red, respectively. The top and bottom five were additionally extracted and displayed by architecture in Tables 3 and 4.

General Performance – Systems B and C, the large 700MHz/p O3900s at ASC and ERDC, the large 1.3GHz/p P4 at NAVO, and the medium-sized 1.3GHz/p P4 at ARSC were consistently top performers, while the small 1.7GHz/p P4s at ARL and ARSC, the medium-sized 375MHz/p P3 at ASC, the small 833MHz/p SC40 at ASC, and the large T3Es at ERDC and AHPCRC were consistently poor performers. For smaller systems with newer architectures, it was not surprising that poor performance was observed, given the premise of the comparison – overall capability. For a larger system of the same type, the ranking would, no doubt, improve. The 400MHz/p X1s at ERDC, AHPCRC, and ARSC exhibited a bi-modal performance with good marks on AERO, Cobalt-60 Standard and Large, and HYCOM Standard, and poor marks on GAMESS Standard and Large, NAMD Standard and Large, and RFCTH Large. For AERO, GAMESS, and NAMD, these results were not surprising given that AERO is a vector code, and GAMESS and NAMD do not vectorize well.

General Performance per Processor – Systems A and B, the small 1.7GHz/p P4s at ARL and ARSC, the small 3.06GHz/p Xeon cluster at ARL, and the 400MHz/p X1s at ERDC, AHPCRC, and ARSC consistently demonstrated good performance density, while the large 375MHz/p P3s at ARL and NAVO, the medium-sized 375MHz/p P3 at MHPCC, and the large T3Es at ERDC and AHPCRC consistently demonstrated poor performance density.

Additional Performance Notables

- System A performed well for both test cases of GAMESS and OOCore, but only moderately well for the other test cases.
- Despite being a generally good performer, System C performed poorly on AERO.
- The small 1.7GHz/p P4s at ARL and ARSC performed moderately well on AERO, despite being generally poor performers.
- The small 3.06GHz/p Xeon cluster at ARL performed poorly on AERO and for both test cases of Cobalt-60, but well for both test cases of GAMESS.
- The medium-sized 400MHz/p O3800s at ARL and ERDC performed well on synthetic tests, but poorly on all other test.
- The small 1GHz/p SC45 at ASC performed poorly for HYCOM Large and OOCore Large, and not much better for the other test cases.
- The large T3E at ERDC performed well for GAMESS Standard, despite being a generally poor performer.
- The medium-sized 375MHz/p P3 at MHPCC performed better than a smaller version with a like architecture at ASC due to its additional size yet still performed poorly.
- The small 1.3GHz/p P4 at MHPCC performed well on AERO and poorly on the synthetic tests.

Additional Performance Density Notables

- System B exhibited a poor performance density (PD) for the synthetic tests despite generally having a good PD.
- System C exhibited a good PD for both test cases of GAMESS and HYCOM Standard, but a poor PD for AERO and the synthetic tests.
- Despite having generally good PDs, the small 1.7GHz/p P4s at ARL and ARSC exhibited relatively poor PDs for the synthetics tests.

- The medium-sized 400MHz/p O3800s at ARL and ERDC exhibited good PDs for the synthetic tests, but poor PDs for everything else.
- The small 833MHz/p SC40 at ASC exhibited a good PD for both test cases of NAMD, but a poor PD for both test cases of Cobalt-60 and OOCore.
- The large 700MHz/p O3900s at ASC and ERDC exhibited good PDs for the synthetic tests, but poor PDs for everything else.
- Despite having generally good PDs, the small 400MHz/p X1s at ERDC, AHPCRC, and ARSC exhibited poor PDs for both test cases of NAMD.
- The medium-sized 833MHz/p SC40 at ERDC exhibited a poor PD for both test cases of OOCore, but a descent PD for both test cases of NAMD and the synthetic tests.
- The small 1.3GHz/p P4 at MHPCC exhibited a good PD for NAMD Large and a poor PD for the synthetics tests.

4. Significance to DoD

The mapping of problems to resources significantly impacts the efficiency of the program, given the diversity of the system architectures and sizes that are available as well as the large span of problems at hand. Providing detailed performance (and performance density) data to users aims at improving this mapping by swaying the users' choice of platforms to those best-suited for their problems.

5. Systems Used

All major unclassified systems within the HPCMP were used.

6. CTA

The computational areas covered by this effort include CCM, CEA, CFD, CSM, and CWO.

Acknowledgements

This mammoth benchmarking effort was conducting by the CS&E Group at ERDC headed by Dr. William A. Ward, Jr., Mr. George Petit at ARL (Raytheon), and Mr. Daniel S. Schornak at ASC (CSC). This work was supported in part by a grant of computer time from the DoD High Performance Computing Modernization Program at the ARL, ASC, ERDC, and NAVO MSRCs.

References

1. Tracy F.T. et al., "A survey of the algorithms in the TI-03 application benchmarking suite with emphasis on linear system solvers." *IEEE Proceedings of the 2003 Users Group Conference*, June 2003, pp. 332–336.

Table 1. Performance Ranking

	AERO	Cobalt60		GAMESS		HYCOM		NAMD		OOCore		RFCTH		Overall	Overall	Overall
	std	std	lg	std	lg	std	lg	std	lg	std	lg	std	lg	Synth	App	Score
System A	6	12	11	4	4	9	8	7	7	3	3	6	6	12	8	8
System B	1	1	1	2	2	2	2	1	1	1	1	1	1	3	1	1
System C	22	2	2	1	1	1	1	2	2	2	2	2	2	4	2	2
U_ARL_IBM_P3_375MHZ_1024P	21	11	12	15	14	15	13	12	14	10	10	13	14	18	12	15
U_ARL_IBM_P4_FED_1.7GHZ_128P	10	25	25	19	18	21	18	17	15	23	19	22	16	25	23	24
U_ARL_LNXI_XEON_3.06GHZ_256P	23	24	24	6	6	12	10	15	20	13	13	19	15	21	15	18
U_ARL_SGI_O3800_400MHZ_512P	17	19	18	21	21	22	21	18	17	21	20	21	18	5	20	14
U_ASC_IBM_P3_375MHZ_516P	25	20	23	20	20	20	19	20	19	17	16	20	23	23	21	22
U_ASC_HP_SC40_833MHZ_56P	19	27	27	26	27	25	27	24	22	25	25	25	25	26	27	27
U_ASC_HP_SC45_1.0GHZ_296P	15	22	19	17	16	19	23	14	13	19	23	17	13	14	19	19
U_ASC_HP_SC45_1.0GHZ_472P	14	14	14	11	10	14	17	9	9	16	18	12	9	11	14	12
U_ASC_SGI_O3900_700MHZ_2032P	9	3	3	5	5	7	5	3	3	5	5	3	3	1	4	3
U_ERDC_CRAY_T3E_600MHZ_1792P	26	16	15	9	15	26	25	22	23	26	26	26	26	17	24	23
U_ERDC_CRAY_X1_400MHZ_60P	7	7	8	27	26	5	12	27	27	12	14	15	24	19	11	13
U_ERDC_HP_SC40_833MHZ_488P	16	17	17	14	12	16	20	11	11	20	22	14	12	13	17	17
U_ERDC_HP_SC45_1.0GHZ_488P	13	13	13	10	9	13	16	8	8	15	17	11	8	10	13	11
U_ERDC_SGI_O3800_400MHZ_504P	18	21	20	22	22	23	22	19	18	22	21	23	19	6	22	16
U_ERDC_SGI_O3900_700MHZ_1008P	11	9	9	8	8	11	11	6	5	11	11	7	7	2	9	5
U_NAVO_IBM_P3_375MHZ_1328P	20	10	10	12	13	10	9	10	10	9	7	10	10	15	10	10
U_Target_IBM_P4_COL_1.3GHZ_1392P	2	4	6	3	3	6	3	4	4	4	4	4	4	7	3	4
U_AHPCRC_CRAY_T3E_600MHZ_1056P	27	23	22	16	19	27	26	23	24	27	27	27	27	22	26	25
U_AHPCRC_CRAY_X1_400MHZ_126P	3	5	4	24	24	3	6	25	25	6	8	8	20	8	6	6
U_ARSC_CRAY_X1_400MHZ_126P	4	6	5	25	25	4	7	26	26	7	9	9	21	9	7	7
U_ARSC_IBM_p690_FED_1.3GHZ_720P	5	8	7	7	7	8	4	5	6	8	6	5	5	16	5	9
U_ARSC_IBM_p690_FED_1.7GHZ_64P	12	26	26	23	23	24	24	21	21	24	24	24	22	27	25	26
U_MHPCC_IBM_P3_COL_375MHZ_736P	24	15	16	18	17	18	15	16	16	14	12	18	17	20	18	21
U_MHPCC_IBM_p690_COL_1.3GHZ_320P	8	18	21	13	11	17	14	13	12	18	15	16	11	24	16	20

Top 5

Bottom 5

Table 2. Performance Per Processor Ranking

	# of p	AERO std	Cobalt60 std	Cobalt60 lg	GAMESS std	GAMESS lg	HYCOM std	HYCOM lg	NAMD std	NAMD lg	OOCore std	OOCore lg	RFCTH std	RFCTH lg	Overall Synth	Overall App	Overall Score
System A	---	6	5	5	3	4	8	8	4	6	5	4	6	6	11	7	5
System B	---	7	4	4	4	5	10	10	5	9	4	5	7	7	17	6	7
System C	---	25	9	15	1	1	5	7	9	13	6	8	15	13	18	9	9
U_ARL_IBM_P3_375MHZ_1024P	1024	23	24	24	24	24	22	24	24	22	18	16	24	24	24	24	24
U_ARL_IBM_P4_FED_1.7GHZ_128P	128	5	7	7	6	6	7	6	2	3	9	9	5	2	16	5	6
U_ARL_LNXI_XEON_3.06GHZ_256P	256	13	15	14	2	2	4	5	6	10	7	7	13	9	14	8	8
U_ARL_SGI_O3800_400MHZ_512P	512	18	21	21	22	22	25	22	18	18	25	24	22	21	6	21	19
U_ASC_IBM_P3_375MHZ_516P	516	21	22	22	20	20	20	19	20	19	16	14	20	22	22	22	22
U_ASC_HP_SC40_833MHZ_56P	56	8	18	17	7	9	12	16	3	2	22	20	8	8	12	11	11
U_ASC_HP_SC45_1.0GHZ_296P	296	12	11	9	13	11	13	13	7	4	14	13	9	11	8	13	10
U_ASC_HP_SC45_1.0GHZ_472P	472	14	12	10	15	14	15	14	11	7	13	18	12	14	9	15	12
U_ASC_SGI_O3900_700MHZ_2032P	2032	20	16	16	19	19	18	20	16	16	21	22	19	19	5	19	16
U_ERDC_CRAY_T3E_600MHZ_1792P	1792	27	27	27	27	27	27	27	27	27	27	27	27	27	27	27	27
U_ERDC_CRAY_X1_400MHZ_60P	60	1	1	1	8	7	1	1	19	21	1	1	1	3	3	1	1
U_ERDC_HP_SC40_833MHZ_488P	488	16	19	18	17	17	17	17	13	12	23	25	17	16	13	17	21
U_ERDC_HP_SC45_1.0GHZ_488P	488	15	13	11	16	15	16	15	12	8	15	19	14	15	10	16	13
U_ERDC_SGI_O3800_400MHZ_504P	504	17	20	20	21	21	24	21	17	17	24	23	21	20	7	20	18
U_ERDC_SGI_O3900_700MHZ_1008P	1008	19	17	19	18	18	19	18	15	15	20	21	18	18	4	18	14
U_NAVO_IBM_P3_375MHZ_1328P	1328	24	25	25	25	25	23	25	25	23	19	17	25	25	25	25	25
U_Target_IBM_P4_COL_1.3GHZ_1392P	1392	9	10	12	14	16	14	12	14	14	11	12	16	17	21	14	20
U_AHPCRC_CRAY_T3E_600MHZ_1056P	1056	26	26	26	26	26	26	26	26	26	26	26	26	26	26	26	26
U_AHPCRC_CRAY_X1_400MHZ_126P	126	2	2	2	11	12	2	2	21	24	2	2	2	4	1	2	2
U_ARSC_CRAY_X1_400MHZ_126P	126	3	3	3	12	13	3	3	22	25	3	3	3	5	2	3	3
U_ARSC_IBM_p690_FED_1.3GHZ_720P	720	10	8	8	10	10	9	9	10	11	10	10	10	12	19	10	15
U_ARSC_IBM_p690_FED_1.7GHZ_64P	64	4	6	6	5	3	6	4	1	1	8	6	4	1	15	4	4
U_MHPCC_IBM_P3_COL_375MHZ_736P	736	22	23	23	23	23	21	23	23	20	17	15	23	23	23	23	23
U_MHPCC_IBM_p690_COL_1.3GHZ_320P	320	11	14	13	9	8	11	11	8	5	12	11	11	10	20	12	17

Top 5
Bottom 5

Table 3. Best and Worst Five Systems by Architecture (Performance).

	AERO std	Cobalt60 std	Cobalt60 lg	GAMESS std	GAMESS lg	HYCOM std	HYCOM lg	NAMD std	NAMD lg	OOCore std	OOCore lg	RFCTH std	RFCTH lg	Overall Synth	Overall App	Overall Score
B	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW	NEW
E		O3900	O3900	O3900	O3900		O3900	O3900	O3900	O3900	O3900	O3900	O3900	O3(8/9)00	O3900	O3900
S	P4-1.3	P4-1.3		P4-1.3	P4-1.3		P4-1.3	P4-1.3	P4-1.3	P4-1.3	P4-1.3	P4-1.3	P4-1.3		P4-1.3	P4-1.3
T	X1	X1	X1			X1										
						O3800						O3800				
W	P3		P3										P3	P3		
O		P4-1.7	P4-1.7	P4-1.7	P4-1.7	P4-1.7	P4-1.7			P4-1.7	P4-1.7	P4-1.7		P4-1.3/1.7	P4-1.7	P4-1.7
R		SC40	SC40	SC40	SC40	SC40	SC40/45	SC40		SC40	SC40/45	SC40	SC40	SC40	SC40	SC40
S	T3E	T3E				T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E		T3E	T3E
T	Xeon Cl	Xeon Cl	Xeon Cl													
				X1	X1			X1	X1				X1			

* The entries in this table are susceptible to system size. Please supplement with data in Table 1.

Table 4. Best and Worst Five Systems by Architecture (Performance Per Processor).

	AERO std	Cobalt60 std	Cobalt60 lg	GAMESS std	GAMESS lg	HYCOM std	HYCOM lg	NAMD std	NAMD lg	OOCore std	OOCore lg	RFCTH std	RFCTH lg	Overall Synth	Overall App	Overall Score
		NEW	NEW	NEW	NEW	NEW		NEW		NEW	NEW					NEW
B														O3900		
E	P4-1.7			P4-1.7	P4-1.7		P4-1.7	P4-1.7	P4-1.3/1.7			P4-1.7	P4-1.7		P4-1.7	P4-1.7
S								SC40	SC40/45							
T	X1	X1	X1			X1	X1			X1	X1	X1	X1	X1	X1	X1
				Xeon Cl	Xeon Cl	Xeon Cl	Xeon Cl									
W	NEW															
O						O3800				O3800	O3800					
R	P3	P3	P3	P3	P3	P3	P3	P3	P3			P3	P3	P3	P3	P3
S										SC40	SC40					
T	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E	T3E
									X1							